

# SmartHBA 2100 and SmartRAID 3100 Software/Firmware Release Notes



# Table of Contents

1. About This Release.....	3
1.1. Release Identification.....	3
1.2. Components and Documents Included in this Release.....	4
1.3. Files Included in this Release.....	4
2. What's New?.....	6
2.1. Features.....	6
2.2. Fixes.....	6
2.3. Limitations.....	20
3. Updating the Controller Firmware.....	26
3.1. Updating the Controller Firmware.....	26
4. Installing the Drivers.....	28
5. Revision History.....	29
Microchip Information.....	30
The Microchip Website.....	30
Product Change Notification Service.....	30
Customer Support.....	30
Microchip Devices Code Protection Feature.....	30
Legal Notice.....	30
Trademarks.....	31
Quality Management System.....	32
Worldwide Sales and Service.....	33

# 1. About This Release

The development release described in this document includes firmware, OS drivers, tools, and host management software for the solutions from Microchip.

## 1.1 Release Identification

The firmware, software, and driver versions for this release are shown in the following table.

**Table 1-1.** Release Summary

<b>Solutions Release</b>	2.7.0
<b>Package Release Date</b>	October 31, 2023
<b>Firmware Version</b>	5.32 B0 <sup>1,2</sup> (basecode e502dcc@HEAD)
<b>UEFI Version</b>	2.2.4
<b>Legacy BIOS</b>	2.2.2
<b>Driver Versions<sup>3</sup></b>	Windows SmartPQI: <ul style="list-style-type: none"> <li>• Windows 2016/2019/2022: 1010.42.0.1020</li> <li>• Windows 10/11: 1010.42.0.1020</li> </ul> Linux SmartPQI: <ul style="list-style-type: none"> <li>• RHEL 7/8/9: 2.1.18-045</li> <li>• SLES 12/15: 2.1.18-045</li> <li>• Ubuntu 16/18/20/21/22: 2.1.18-045</li> <li>• Debian 10/11: 2.1.18-045</li> <li>• CentOS 7/8: 2.1.18-045</li> <li>• Oracle Linux 7/8: 2.1.18-045</li> <li>• Citrix XenServer 8: 2.1.18-045</li> </ul> VMware SmartPQI: <ul style="list-style-type: none"> <li>• VMware 7.0: 4330.0.116</li> </ul> FreeBSD/Solaris SmartPQI: <ul style="list-style-type: none"> <li>• FreeBSD 12/13: 4280.0.1007</li> <li>• Solaris 11: 11.4120.0.1005</li> </ul>
<b>Management Software (arccconf, maxView™, Event Monitor, BootUSB)</b>	B25339
<b>PLDM</b>	6.10.14.0

### Notes:

1. Downgrading to 1.04 B0 or older builds from this release or prior 1.29 releases may cause the board to not boot or have supercap errors due to an incompatibility in SEEPROMs between this release and prior releases. Refer to the section "[3. Updating the Controller Firmware](#)" to downgrade an existing board.
2. If the firmware running on the board is older than 0.01 B594, existing data in the logical drives must be backed up if it needs to be used after the upgrade. After the upgrade from firmware prior to 0.01 B594, the logical drives will need to be recreated.
3. Only run the driver on firmware 0.01 build 500 or later.
4. Windows 11 Inbox and OOB driver is supported.

## 1.2 Components and Documents Included in this Release

Download the firmware, drivers, host management software, and supporting documentation for your SmartHBA 2100/SmartRAID 3100 controller and SmartRAID 3100 and SmartRAID 3100 controller solutions from the Microchip Web site at <https://start.adaptec.com>

## 1.3 Files Included in this Release

This release consists of the files listed in the following tables:

### Firmware Files

**Table 1-2.** Firmware Files

Component	Description	Pre-Assembly Use	Post-Assembly Use
SmartFWx100.bin	Programmable NOR Flash File Use to program NOR Flash for boards that are already running firmware.	—	X
SmartFWx100.fup	Programmable NOR Flash File Used for PLDM type 5 firmware flashing for boards that are already running firmware.	—	X

**Table 1-3.** Firmware Programming Tools

Tool	Description	Executable
Arcconf romupdate	The command allows to upgrade/downgrade the firmware and BIOS image to the controller.	Refer to <a href="#">Table 1-7</a>
maxView™ firmware upgrade wizard	The firmware upgrade wizard allows to upgrade/downgrade the firmware and BIOS image to one or more controller(s) of same model in the system.	Refer to <a href="#">Table 1-7</a>

### Driver Files

**Table 1-4.** Windows Storport Miniport SmartPQI Drivers

Drivers	Binary	Version
Server 2022, Server 2019 and Server 2016 Windows 10 (version 21H2) and 11	SmartPqi.sys	x64
	SmartPqi.inf	x64
	Smartpqi.cat	x64

**Table 1-5.** Linux SmartPQI Drivers

Drivers	Version
Red Hat Enterprise Linux 9.0 <sup>2</sup> , 8.6 <sup>2</sup> , 8.5, 8.4, 7.9, 7.8	x64
CentOS 8.4, 8.3, 8.2, 8.0, 7.9, 7.8	x64
SuSE Linux Enterprise Server 12 <sup>1</sup> , SP5, SP4	x64
SuSE Linux Enterprise Server 15 SP4 <sup>2</sup> , SP3, SP2	x64
Oracle Linux 7.9 UEK6U3	x64
Oracle Linux 8.5, 8.4 UEK6U3	x64
Ubuntu 22.04, 21.04	x64
Ubuntu 20.04.4, 20.04.3, 20.04	x64
Ubuntu 18.04.5, 18.04.4, 18.04	x64

.....continued	
Drivers	Version
Ubuntu 16.04.5	x64
Debian 11.1, 10.10, 10.05	x64
Citrix xenServer 8.2, 8.1, 8.0	x64
Fedora 35 (inbox only)	x64

**Notes:**

1. To mitigate against the Spectre Variant 2 vulnerability, the RHEL 6u9/RHEL7u4/RHEL7u5 and SLES11 SP3 and higher drivers have been compiled to avoid the usage of indirect jumps. This method is known as "Retpoline".
2. New OS support—minimally tested drivers in this release. Fully supported drivers are expected in the next release.

**Table 1-6.** FreeBSD, Solaris, and VMware SmartPQI Drivers

Drivers	Version
FreeBSD 13, 12.3	x64
Solaris 11.4	x64
VMware 7.0 U3/U2	x64

**Host Management Software****Table 1-7.** Host Management Utilities

Description	OS	Executable
ARCCONF Command Line Utility	Windows x64 Linux x64 VMware 6.5 and above XenServer FreeBSD x64 Solaris x86	See the Arccnf download package for the OS-applicable installation executable.
ARCCONF for UEFI		Included as part of the firmware downloadable image.
maxView™ Storage Manager	Windows x64 Linux x64 VMware EXSi 6.5 and above XenServer	See the maxView Storage Manager download package for the OS-applicable installation executable.
maxView™ vSphere Plugin	VMware 6.5 and above	See the VMware maxView Storage Manager download package for the OS-applicable installation executable.
Boot USB (offline or pre-boot) for ARCCONF and maxView Storage Manager	Linux x64	See the maxView BootUSB download package for the .iso file.

## 2. What's New?

This section shows what's new in this release.



**Important:** Updated maxView to address log4j vulnerabilities.

### 2.1 Features

The following table lists features supported for this release.

**Table 2-1.** Feature Summary

Feature		Supported in this Release	Future Release
UEFI Driver, Boot Support		X	
Legacy Boot Support		X	
Dynamic Power Management		X	
SMR Drive Support	Enumeration, Unrestricted Command Flow-Through	X	
	SATL Translation for HA/HM SMR Management	X	
	Identify All Drive Types	X	
Driver Support	Windows	X	
	Linux	X	
	VMware	X	
	FreeBSD	X	
	Solaris	X	
	OS certification	X	
Out of Band interface selection support of MCTP or PBSI		X	
Flash Support		X	
MCTP BMC Management		X	
SED Local Key Management		X	
Configurable Big Block Cache Bypass		X	
Green Backup Support for SmartRAID		X	
4Kn Support in RAID		X	

### 2.2 Fixes

#### 2.2.1 Firmware Fixes

##### 2.2.1.1 Fixes and Enhancements for Firmware Release 5.32 B0

This release includes the following fixes and enhancements:

- Added support for Managed SED Local mode.
- Added support for Multi-Actuator disk drives.
- Added support for reporting previous cache read/write setting to PLDM when only\_read cache setting happens.

- Added support for random logical device deletion without backup power source.
- Fixed an issue where the encrypted data is not accessible for a RAID 50/60 logical drive when it was failed and healed using "Heal Array".
  - Root Cause: During the "Heal Array" operation on a failed encrypted logical drive, firmware creates a new Data Encryption Key (DEK) for the healed encrypted logical drive. This causes the data to be inaccessible because the new DEK is different from the original DEK used to encrypt the data for those physical drives.
  - Fix: During a "Heal Array" operation, firmware keeps the previously used DEK so the data is accessible.
  - Risk: Low
- Fixed an issue where creating logical drive failed with pre-existing maxCache but no backup power source.
  - Root Cause: Firmware is reporting different non-battery-backed cache memory sizes for different states of the backup power source. This results in the host management software sending the wrong cache configuration sizes when creating a new logical drive.
  - Fix: Non-battery-backed cache memory size is made uniform irrespective of the Backup Power Source state so the host management software can send correct cache configuration sizes.
  - Risk: Low
- Fixed an issue where controller hangs under heavy I/O during rebuild of any logical drive except RAID 6.
  - Root Cause: One of the firmware threads acquired a lock for a larger range of LBAs to perform a rebuild and waited for data transfer resources. Simultaneously, one of the other firmware threads acquired all available data transfer resources to service host I/O for an already locked range of LBAs. This resulted in a firmware deadlock.
  - Fix: Implemented a partial rebuild method to minimize the required data transfer resources. Also, allocated enough data transfer resources to run a minimum of one partial rebuild iteration during heavy I/O.
  - Risk: Medium
- Fixed an issue where the state change of the logical drive is not appropriate when marking a failed physical drive.
  - Root Cause: A logical drive with a spare drive is created and when a drive fails, the state of logical drive changes to DEGRADED state and then to NEEDS\_REBUILD state. Firmware starts the rebuild on the spare drive. Firmware then handles the same drive failure a second time and the state of the logical drive changes to DEGRADED again. Consequently, the full process of rebuild starts again.
  - Fix: Added an extra parameter in the firmware to reflect the drive failure status. When the firmware fails the drive again, based on the drive failure status, the state of the logical drive does not change again.
  - Risk: Low
- Fixed an issue where the slot number shows 'unknown' for direct attached drives.
  - Root Cause: The firmware fails to fill the bay information properly for direct attached drives if it is not connected in an enclosure. During boot time, if the SGPIO backplane configuration is based on the host-provided configuration, then the firmware has to get the host-provided configuration bit stream and retrieve the bay information from it. The firmware failed to fill the bay and box information as SGPIO data READ did not happen properly, hence firmware uses the default configuration. Later, when the firmware tried to update the box and bay information into the default configuration for the direct attached drives, drives

- that were not connected in an enclosure were skipped and the bay number remained as DEVICE\_NOT\_IN\_BAY even if the drive is present and exposed to the host.
- Fix: Remove initialization of the bay information from host-provided configuration bit stream, so the firmware can assign a logical PHY number for the direct attached drive's bay.
  - Risk: Medium
- Fixed an issue where second transformation was suspended after an unsafe system reboot during first transformation.
    - Root Cause: When the transformation method is switched from internal cache-based to disk-based method, the firmware did not clear the internal cache status. When the transformation ended for the current volume and started for the next volume, the firmware suspended the transformation due to uncleared internal cache status. If the backup power is removed, the internal cache status will never Reset and the transformation will freeze.
    - Fix: Once the transformation backup method is changed safely from internal cache-based method to disk-based method, the internal cache status is cleared.
    - Risk: Low
  - Fixed an issue where the persistent event logs were not captured in the debug log.
    - Root Cause: Firmware uses a circular buffer mechanism for storing events in the memory. If the host is consuming the firmware event buffer, then the firmware needs to increase the head pointer each time with the help of event size. When the head pointer reaches an end but only a few bytes are left to read which is less than one event size, the firmware needs to ignore it. Here, the firmware is not ignoring the event and finally it returns as an empty buffer to the host.
    - Fix: Calculate the head pointer and check against the end of persistent event logging. If leftover bytes are less than one event size, then ignore them.
    - Risk: Low
  - Fixed an issue where the drive slot number shown for a failed drive is wrong.
    - Root Cause: If the SES enclosure has a valid additional status page, then the bay number for each drive in the enclosure is used by the firmware from the additional status page information. If the drive has an invalid SAS address in the additional status page, the bay number is assigned from the expander drive slot number by the firmware. When the firmware is traversing through physical drives to validate the additional status page, the validation flag is overwritten with the latest drive's additional status page validity status. This overwriting resulted in firmware not updating the bay number if the drive had an invalid SAS address in the additional status page.
    - Fix: Stop overwriting the validation flag of the drive's additional status page when any drive in the list has an invalid status.
    - Risk: Low
  - Fixed an issue where logical drives may be incorrectly failed after a rebuild.
    - Root Cause: An array consisting of multiple logical drives and multiple spare drives will undergo a rebuild when one of the physical drives has failed. Each logical drive is sequentially rebuilt using a spare drive. If some of the logical drives have completed the rebuild and a second physical drive fails, then firmware incorrectly reports the previously rebuilt logical drives as failed instead of degraded and the logical drive that is currently being rebuilt is correctly marked as failed. The firmware incorrectly failed the previously rebuilt logical drives because it removed the spare drives and the failed drive from those logical drives.
    - Fix: Firmware keeps the spare drives as part of the rebuilt logical drives so they will correctly report a degraded state.



- Risk: High
- Fixed an issue where the MCTP Discovery Notify setting cannot be configured.
  - Root Cause: SEEPROM has one bit reserved to enable or disable the MCTP discovery notification. This bit is not checked in the firmware while initializing the MCTP driver during boot. Regardless of this SEEPROM bit value, MCTP discovery notification is kept enabled. Therefore, the command given by host management software to enable or disable MCTP Discovery Notify is not working.
  - Fix: Check the MCTP discovery notification bit from SEEPROM before enabling it. Also, ensure the Firmware is used with arconf/Maxview B25335 version or later.
  - Risk: Low
- Fixed an issue where firmware does not abort all pending requests when the device LUN reset occurs.
  - Root Cause: Upon device LUN Reset, firmware does not abort requests pending on device's internal queue.
  - Fix: Upon device LUN Reset, firmware aborts all requests pending on device's internal queue.
  - Risk: Low
- Fixed an issue where the PBSI multi-LUN field remains unset for multi-LUN drives.
  - Root Cause: PBSI support is not available for multi-LUN devices.
  - Fix: Added PBSI support for multi-LUN devices.
  - Risk: Low
- Fixed an ~8% performance drop on HBA mode for 4K/8K random read with the I/O queue depth of 64 or higher after upgrading the firmware version from v1.32 to v4.11.
  - Root Cause: The default SAS drive queue depth was incorrectly set to 32.
  - Fix: Set the default SAS drive queue depth back to 64.
  - Risk: Low
- Fixed an issue where firmware continuously retries a command to the SEP device that was completed with the TASK SET FULL status.
  - Root Cause: Firmware retries indefinitely on any command to the SEP device that was completed with TASK SET FULL status and the request's retry count was still greater than 1.
  - Fix: Firmware decreases the request's retry count for all the retry commands to the SEP device regardless of the SCSI error status.
  - Risk: Low
- Fixed an issue where the encrypted single drive RAID 0 logical drive was reported with state NEED\_CRYPTO\_KEY after system reboot.
  - Root Cause: When the physical drive was replaced and the encrypted single drive RAID 0 logical drive was re-enabled, the key that is used to encrypt the data was not written back to the logical drive media. The key was used from within the controller's internal memory. When the system was rebooted, the key is no longer in the controller memory and the key is not on the logical drive either, so it is not present in the current system.
  - Fix: Write the key back to the media of all physical drives contained in the logical drive when the logical drive is re-enabled.
  - Risk: Low
- Fixed an issue with a UBM backplane where not all drives were discovered.

- Root Cause: Backplane supported 10 DFC connectors and firmware assumed no more than eight.
  - Fix: Increased max DFC connectors to 15.
  - Risk: Low
- Fixed an issue where the product ID of an enclosure was not showing correctly.
  - Root Cause: Product IDs that had a space in the middle of the string were being truncated.
  - Fix: Corrected the code to ensure any characters after the space are displayed.
  - Risk: Low
- Fixed an issue where the Predictive Fail (PF) LED did not blink during predictive spare rebuild.
  - Root Cause: When rebuild started, only a call to blink the target rebuilding drive was called.
  - Fix: When rebuild starts, update LED states for all drives instead of just the rebuilding drive. PF LED blinks only if logical drive is Fault tolerant.
  - Risk: Low
- Fixed a possible controller hang during a clear configuration operation if the logical drive did not finish RPI and has a physical drive with index D254.
  - Root Cause: The drive with index D254 did not have its RPI information initialized that caused invalid RPI state information to be used by firmware. The invalid RPI state told firmware RPI was still running, which meant the clear configuration command is stuck waiting for RPI to halt.
  - Fix: Ensure RPI information is initialized for drives assigned to D254.
  - Risk: Low
- Fixed long SATA SSD TRIM causing hang.
  - Root Cause: While creating a logical drive with SATA SSDs, if an Over-Provisioning Optimization (OPO) operation is performed on SSDs, no host IOs are allowed to complete. This makes firmware assume there is a hang and firmware may trigger a lockup.
  - Fix: Resolved the lockup by considering OPO operations as normal I/O. Also, to speed up the OPO operation for SATA SSDs, increase the number of bytes being trimmed per operation to 512 MiB, which has been observed to reduce the OPO operation time to 22 seconds per terabyte.
  - Risk: Medium
- Fixed continuous Predictive Failure LED blinking during an ongoing Predictive Spare Rebuild.
  - Root Cause: When the drive is hot removed the `update_led` function is only called for the drive which was hot removed.
  - Fix: Update LED behavior for all drives when the drive is hot removed. This ensures the Predictive Failure LED is turned off during a rebuild operation.
  - Risk: Low
- Fixed a potential UBM backplane Flash issue where image start index is not aligned with the start sector.
  - Root Cause: Controller was not calculating start index, and assumed the start index aligns with the start sector.
  - Fix: Calculate the start index correctly.
  - Risk: Low
- Fixed an issue where the Real Time Clock (RTC) timestamp was not sent to the SES based storage enclosure SEPs attached to internal connectors of the controller.

- Root Cause: Firmware was only sending the RTC timestamp to SEPs that were connected to external connectors of the controller.
- Fix: Removed check for externally connected enclosure so timestamp is sent to all SEPs.
- Risk: Low
- Fixed an issue with two UBM backplanes connected to one connector and one UBM backplane incorrectly identified itself as not on a bifurcated cable.
  - Root Cause: The logic for each connector was trying to take ownership of the drives that resulted in controller locking up because drives should only be owned by one box.
  - Fix: Do not discover drives on second connector if drives have already been marked.
  - Risk: Low
- Fixed an issue with sending a PLDM event when controller password is entered and no encrypted logical drives are present.
  - Root Cause: Firmware was assuming an encrypted logical drive was available to send an event when the controller password is entered.
  - Fix: Send an event even if no encrypted logical drive is present.
  - Risk: Low

## 2.2.2 UEFI Fixes

**Note:** Microsoft signed and secure boot is supported.

### 2.2.2.1 Fixes and Enhancements for UEFI Driver 2.2.4/Legacy BIOS 2.2.2

This release includes the following UEFI fixes and enhancements:

- Added a new HII menu that will attempt to re-enable a previously failed logical drive whose physical drives are back online later.
- Driver health error codes are consolidated from 0x17xx and 0x18xx series to a 0x19xx series.
- Added an Unlock Controller option in the HII menu when controller password is set for Controller-Based Encryption.
- Added new HII options to enable and configure controller managed SED based encryption for physical and logical drives.
- Fixed a controller lockup during transition from preboot to OS.
  - Root Cause: The UEFI driver sends I/O commands to the controller without using a timeout value. During the transition from preboot to an OS, if an I/O command is still pending and the UEFI driver triggers a command interface change, the controller encounters a lockup.
  - Fix: Timeout value added for I/O commands.
  - Risk: Low
- Fixed an issue of incorrect location representation for non-disk devices such as SES and expander devices.
  - Root Cause: Location represented as `port:box:bay` for non-disk devices. Bay number is not applicable for non-disk devices.
  - Fix: Location for non-disk devices updated to only show `port:box` format.
  - Risk: Low

## 2.2.3 Driver Fixes

### 2.2.3.1 Fixes and Enhancements for Linux Driver Build 2.1.18-045

This release includes the following fixes and enhancements.

- Added support for Multi-Actuator disk drives.
- Added support for displaying controller firmware version in the OS message log. The controller firmware version is printed to OS message log during driver initialization.
- Added support for a controller ready timeout module parameter (`ctrl_ready_timeout`). The valid range is 0 or 30–1800 seconds. The default value is 0, which causes the driver to use a timeout of 180 seconds (3 minutes).
- Added support for deleting a LUN via sysfs using the following syntax:

```
echo 1 > /sys/block/sdX/device/delete
```

- Added module parameter to disable managed interrupts (`disable_managed_interrupts=1`).
- Fixed a race condition where the driver can access the RAID map when using IOBypass during a RAID configuration change.
  - Root Cause: A race condition in the driver might cause it to access a stale RAID map when a logical drive is reconfigured.
  - Fix: Modified the driver logic to
    - Invalidate a RAID map at an early stage when a RAID configuration change is detected
    - Switch to a new RAID map only after the driver detects that the RAID map has changed
  - Risk: Low
- Fixed an issue where the `sg_map` tool issues `SCSI READ BLOCK LIMITS (0x5)` command and the firmware never completes it, causing a system call trace and `sg_map` hang.
  - Root Cause: Driver is sending an incorrect data direction flag for the RAID path request.
  - Fix: Corrected the data direction flag for the RAID path request.
  - Risk: Low
- Fixed an issue where PQI Reset might fail with an error “– 6” if firmware takes more than 100 ms to complete Reset.
  - Root Cause: Method used by the driver to detect controller firmware crash during PQI Reset was incorrect in some cases.
  - Fix: Changed method used by the driver to detect controller firmware crash during PQI Reset.
  - Risk: Low

### 2.2.3.2 Fixes and Enhancements for FreeBSD Driver Build 4280.0.1007

This release includes the following enhancements and fixes:

- Fixed an issue where logical drives created on the hard disk drives were not listed with the “`geom disk list`” command.
  - Root Cause: The `bus_dmamap_sync` operation which is performed before device access was being handled incorrectly.
  - Fix: Handle the `bus_dmamap_sync` operation correctly.
  - Risk: Low
- Fixed an issue of the expanded capacity of a logical drive not reflecting in the OS after modifying the size using `Arcconf`.
  - Root Cause: Current FreeBSD code detects a capacity expansion by checking the logical drive status. Firmware may cycle through the logical drive states (OK, NEEDS\_EXPAND, EXPANDING, OK) faster than the driver can send commands to check for the logical drive state changes. This results in the driver seeing the logical drive as OK and failing to detect the logical drive expansion.

- Fix: Driver detects the event type `PQI_EVENT_TYPE_LOGICAL_DEVICE (0x5)` due to logical drive expansion and sets the rescan flag. Later, it triggers reprobe of all logical drives that includes sending `SCSI READ CAPACITY` command and updating the size that is visible to the OS.
- Risk: Low
- Fixed an issue where OS commands are getting hung after executing on logical drives that are created using SSD.
  - Root Cause: Driver was not properly handling the data direction flag for IOBypass.
  - Fix: Corrected the data direction flag for IOBypass request.
  - Risk: Low
- Fixed an issue where the `SCSI READ BLOCK LIMITS (0x5)` command is never completed by firmware and a `TMF ABORT` is observed.
  - Root Cause: Driver is sending an incorrect data direction flag for the RAID path request.
  - Fix: Corrected the data direction flag for the RAID path request.
  - Risk: Low

### 2.2.3.3 Fixes and Enhancements for Solaris Driver Build 11.4120.0.1005

There are no known fixes for this release.

### 2.2.3.4 Fixes and Enhancements for Windows Driver Build 1010.42.0.1020

- Added support for Multi-Actuator drives.
- Added driver internal ring buffer logging that allows the driver to log important messages to a driver allocated ring buffer memory.
- Fixed an issue where driver accesses Command Descriptor Block's (CDB) NULL pointer and BSOD occurs.
  - Root Cause: The SmartPQI driver parses the SRB with CDB length = 0 and sets the CDB pointer to NULL. When the driver accesses the NULL CDB pointer an invalid memory access occurs.
  - Fix: The build SCSI request in the driver's SRB routine will return invalid command status when the CDB pointer is NULL.
  - Risk: Low
- Fixed an issue where driver might fail to load intermittently after a dirty system shutdown.
  - Root Cause: There are two problems that can occur after a dirty system shutdown.
    - The first is that the driver can get confused about which mode the controller is in and request an unnecessary soft reset.
    - The second problem is that after requesting the unnecessary soft reset, the driver can access controller registers prematurely before the controller has completed the soft reset. This can result in the driver misinterpreting the state of the controller firmware.
  - Fix: Added additional checks to the driver initialization logic to make its Controller mode detection more robust and to prevent misinterpreting controller registers when the controller firmware is not fully up and running.
  - Risk: Low
- Fixed a BSOD due to driver and controller hardware not supporting greater than 16 byte CDB.
  - Root Cause: BSOD is caused by driver copying greater than 16 bytes to the IOBypass request, thus overwriting critical error index field. Therefore, the IOBypass error index field used is out of range and causes the driver to access a bad address.
  - Fix: Added check in build I/O command path to check for CDBs greater than 16 byte and if found, reject the command at the driver level as an invalid command or a command that is not supported.

- Risk: Low

### 2.2.3.5 Fixes and Enhancements for VMware Driver Build 4330.0.116

This release includes the following enhancements and fixes:

- Added support for Multi-Actuator disk drives.
- Added support for multiple tag table to improve performance by optimizing tag and requesting structure allocation.
- Added support for `ScsiAdapterCheckTarget` which performs a device lookup and returns true only when specified adapter/channel/target exists and is exposed as a SCSI device.
- Added a module parameter `CtrlReadyTimeoutSecs` for controller ready timeout. The valid range is 30–1800 seconds. The default value is 120 seconds.
- Fixed an issue where the `SCSI READ BLOCK LIMITS (0x5)` command is never completed by firmware and a `TMF ABORT` is issued in ESXi.
  - Root Cause: Driver is sending an incorrect data direction flag for the RAID path request.
  - Fix: Corrected the data direction flag for the RAID path request.
  - Risk: Low
- Fixed an issue where during driver initialization a warning message appears about the DMA alignment setting, instead of an informational message.
  - Root Cause: Improper flag on the message.
  - Fix: Change message level `WARN` to `INFO`.
  - Risk: Low
- Fixed an issue where PSOD is observed during driver load.
  - Root Cause: Driver was creating a greater number of completion worlds than allowed, which resulted in a SCSI layer deadlock.
  - Fix: Set the SCSI completion world to the maximum supported value.
  - Risk: Medium
- Fixed an issue with cards that are running on older firmware where firmware features previously were potentially disabled due to incorrect placement in the firmware features table.
  - Root Cause: Firmware feature parsing logic on the driver's side was incorrectly skipping some features that must have been enabled regardless of firmware's maximum known features. Firmware does not have the maximum known feature bit set, and the code stops evaluating the feature list when it finds this issue.
  - Fix: Lean on already implemented method of traversing valid feature entries using the `num_elements` field in the PQI firmware feature table, and reject support for any bit positions outside the valid bytes indicated by firmware.
  - Risk: Low
- Fixed an issue where the driver reports an error when the unsupported `SCSI Maintenance IN (0xA3)` command with service action "report supported opcode" (`0xC`) is sent to the logical drive.
  - Root Cause: Firmware does not support command `0xA3` with service action `0xC` and it returns an error to the driver.
  - Fix: Suppress error logs reported for `SCSI Maintenance IN (0xA3)` with service action `0xC`.
  - Risk: Low

## 2.2.4 Management Software Fixes

### 2.2.4.1 Fixes and Enhancements for Arconf/maxView Build 4.11.00.25339

Microchip strongly recommends the maxView users to update to the latest version of the tools to avoid a security vulnerability that has since been resolved.

### 2.2.4.2 Fixes and Enhancements for Arconf/maxView Build B25335

This release includes the following fixes and enhancements for arconf/maxView:

- Removed the Log4J library usage completely from maxView.  
**Note:** Microchip strongly recommends users of maxView update to the latest version of the tool to avoid the security vulnerabilities with the previous releases.
- Added support for managed SED in arconf/maxView.
- Added TASK and GETSTATUS commands support in UEFI Arconf.
- Added ASIC minor version display in maxView and Arconf.
- Added support in maxView and Arconf to sort the controllers based on the bus number, if the slot ID is not valid.
- Fixed an issue where maxView create logical drive selection check box was invisible in older version of the Edge browsers..
  - Root Cause: maxView create logical drive selection table had alignment issues in older version of the Edge browser when the column width was specified in percentage.
  - Fix: The column width of the table is changed to pixel instead of percentage to resolve the alignment issue in the older edge browser.
  - Risk: Low
- Fixed an issue where maxView was failing to modify the maxCache logical device's cache policy to write-through in ESXi 7.xx.
  - Root Cause: ESXi Redfish server had issues in accepting the HTTP PATCH request from maxView.
  - Fix: Updated the maxView HTTP PATCH request call to comply with ESXi redfish server.
  - Risk: Low
- Fixed an issue where Arconf was not displaying the controller PHY error log information.
  - Root Cause: PHY error log was disabled in Arconf.
  - Fix: Added changes to enable the PHY error log from Arconf.
  - Risk: Low
- Fixed an issue where Arconf results in segmentation Fault when collecting the support archive.
  - Root Cause: Arconf resulted in segmentation Fault when retrieving invalid vendor specific diagnostic page from the drive.
  - Fix: Added changes in Arconf to skip reading the invalid vendor specific diagnostic page.
  - Risk: Low
- Fixed an issue where Arconf was not displaying the updated expander firmware version after upgrading the expander firmware.
  - Root Cause: The controller is returning the older expander firmware version instead of recently updated expander firmware version until the next controller power cycle.
  - Fix: Added changes in Arconf to send the SCSI inquiry command to the expander for retrieving the updated expander firmware version.
  - Risk: Low
- Fixed an issue of I<sup>2</sup>C devices not being detected when "SMBUSCHANNEL" is set to "ENABLE".

- Root Cause: Arcconf disabled the SMBUSCHANNEL when user enabled it.
  - Fix: Added changes in Arcconf to set the correct user input value for SMBUSCHANNEL.
  - Risk: Low
- Fixed an issue where Arcconf fails to set ATAPASSWORD for a SATA device.
  - Root Cause: Arcconf ATAPASSWORD command was disabled.
  - Fix: Added changes in Arcconf to enable the ATAPASSWORD command.
  - Risk: Low
- Fixed an issue where Arcconf did not allow changing the write cache setting to write-through for RAID 5 maxCache logical device.
  - Root Cause: Write-through setting for maxCache write cache setting was aborted due to invalid condition check.
  - Fix: Added changes to remove the invalid condition check and allow write-through setting for maxCache logical device.
  - Risk: Low
- Fixed an issue where in UEFI Arcconf where remove hot spare command was failing when array ID was specified.
  - Root Cause: Parsing of array ID in the UEFI Arcconf remove hot spare command was not proper.
  - Fix: Added changes to fix the array ID parsing issue.
  - Risk: Low
- Fixed an issue in Arcconf where an invalid SMART attribute was displayed in the drive SMART STATS.
  - Root Cause: Arcconf was retrieving invalid attributes beyond the SMART attributes range.
  - Fix: Added changes in Arcconf to block the retrieval of invalid attributes beyond the SMART attributes range.
  - Risk: Low
- Fixed an issue where Arcconf was not displaying the connector information for the enclosure device.
  - Root Cause: Connector ID for the enclosure device was missing.
  - Fix: Added changes to display the connector ID for the enclosure device in Arcconf.
  - Risk: Low
- Fixed an issue where maxView was displaying invalid supercap alert message for the controller that does not support Supercap.
  - Root Cause: Supercap alert message was displayed for the controller that does not support Supercap.
  - Fix: Added changes to display the Supercap alert message only for controllers that has Supercap support.
  - Risk: Low
- Fixed an issue where Arcconf failed to execute SLOTCONFIG sub-command.
  - Root Cause: Arcconf SLOTCONFIG sub-command was disabled.
  - Fix: Added changes to enable SLOTCONFIG sub-command from Arcconf.
  - Risk: Low
- Fixed an issue where UEFI Arcconf was not displaying array information for physical device.
  - Root Cause: Array property on the physical device was not implemented in the UEFI Arcconf



- Fix: Added changes to display the array property in GETCONFIG command for the physical device.
- Risk: Low
- Fixed an issue where Arcconf was not displaying the negative temperatures for certain drive models.
  - Root Cause: Arcconf failed to decode the negative temperature for the specific drives.
  - Fix: Added changes to display the correct negative temperature for the supported drives.
  - Risk: Low

### 2.2.4.3 Fixes and Enhancements for PLDM Release 6.10.14.0

This release includes the following fixes and enhancements:

- Added support for the following Redfish ACTION requests:
  - Drive.#SecureErase
  - Drive.#Reset
  - Storage.#ResetToDefaults
    - Note:** ResetToDefaults does not support a ResetType of ResetAll when encrypted volumes exist on the controller. The user must first either delete or decrypt any encrypted volumes prior to issuing such an ACTION request.
- Added support for Redfish PATCH requests for the following properties:
  - Volume.DisplayName
  - Volume.Links.DedicatedSpareDrives
  - Volume.IOPerfModeEnabled
  - Volume.ReadCachePolicy
  - Volume.WriteCachePolicy
  - Drive.LocationIndicatorActive
  - Drive.WriteCacheEnabled
    - Note:** This PATCH is unsupported for drives configured as a Volume's data drive connected to a controller.
  - StorageController.ControllerRates.ConsistencyCheckRatePercent
  - StorageController.ControllerRates.RebuildRatePercent
  - StorageController.ControllerRates.TransformationRatePercent
- Added the following Redfish alerts:
  - DriveOffline
  - DriveMissing
  - DriveMissingCleared
  - DriveOfflineCleared
  - VolumeOffline
  - VolumeOfflineCleared
  - BatteryMissing
  - BatteryFailure
  - BatteryCharging
  - BatteryOK
  - ControllerDegraded
  - ControllerFailure

- ControllerPreviousFailure
- ControllerPasswordRequired
- ControllerPasswordEntered
- Added the following property to the Redfish GET response for VolumeCollection resources:
  - VolumeCollection.@Redfish.CollectionCapabilities.MaxMembers
- Added support for firmware updates for physical drives.
- Added a new OperationName value of 'Reverting' which is used for Redfish GET responses targeting self-encrypting drives undergoing a revert.
- Added a new EncryptionType of 'NativeDriveEncryption' for logical drives which are secured using SED-based encryption.
- Fixed an issue where events are sent continuously if the host does not respond to PlatformEventMessage.
  - Root Cause: For an asynchronous event receiver, there was no logic to cap the number of retries when an event subscriber never responds to a PlatformEventMessage request.
  - Fix: Added logic to cap the number of retries to three. After reaching the maximum limit of retries, no further events will be sent until new events are pushed in or event subscriber resets the event queue.
  - Risk: Medium
- Fixed an issue where Redfish CREATE requests for logical drives fail while using 4 Kn drives as the data drives.
  - Root Cause: The API, which creates a logical drive with 4 Kn drives returned an error causing the CREATE operation to fail.
  - Fix: Updated the API to create a logical drive on 4 Kn drives correctly.
  - Risk: Low
- Fixed an issue where Storage.Status.HealthRollup erroneously reports a value of Warning instead of OK when a logical drive has an InitializationMethod of Foreground.
  - Root Cause: The calculation of the controller's health roll-up considers a logical drive undergoing foreground initialization to have a health value of 'Warning' as it is unavailable to the host OS during that time.
  - Fix: Revised the calculation of the controller health roll-up to publish a value of 'OK' if a child logical drive is queued for or undergoing foreground initialization and no other factors will cause a health status other than 'OK'.
  - Risk: Low
- Fixed an issue where controller firmware update progress is underestimated.
  - Root Cause: The percent complete increment was not adjusted for cases where two flashes are required.
  - Fix: The number of flashes required is now checked when initializing the percent complete increment.
  - Risk: Low
- Fixed an issue where StorageController.CacheSummary.Status.State will have a value of Enabled on controllers lacking a cache module.
  - Root Cause: The mechanism for determining if a cache module was attached was not checking the right fields in the right conditions.
  - Fix: Changed the code to correctly determine the presence of a cache module and correctly report it. During the investigation of this issue it was also determined that other fields in the cache summary also needed to be changed.

If no cache module is attached, the following fields should be as shown here:

- TotalCacheSizeMiB: 0
- PersistentCacheSizeMiB: 0
- Status.Health: OK
- Status.State: Absent

If the controller has no cache module but supports RAID modes, the following fields should be as shown here:

- TotalCacheSizeMiB: The total memory size of the controller
- PersistentCacheSizeMiB: 0
- Status.Health: Warning
- Status.State: Disabled

- Risk: Low
- Fixed an issue where a Port's ServiceLabel did not contain the parent StorageController's slot number.
  - Root Cause: The Port READ response function was duplicating the Port Name property while publishing the ServiceLabel property.
  - Fix: Revised the ServiceLabel value string to contain both the controller slot number and port name in the format "Slot=x:Port=y".
  - Risk: Low
- Fixed an issue where an incorrect completion code is sent for Redfish requests which encounter an error.
  - Root Cause: RDEOperationInit requests which encounters an error returns ERROR\_OPERATION\_FAILED. This does not conform to the DMTF PLDM Type 6 spec DSP0218.
  - Fix: Modified the processing of RDEOperationInit requests to conform to the DMTF PLDM Type 6 specification DSP0218, where any Type 6 operation which encounters an error and responds with extended error information must return ERROR\_UNSUPPORTED instead of ERROR\_OPERATION\_FAILED.
  - Risk: Low
- Fixed an issue where an incorrect strip size is applied while creating logical drive's resources on 4 Kn drives.
  - Root Cause: The calculation of the strip size in blocks used when creating the logical drive assumed a physical drive block size of 512 bytes.
  - Fix: Modified the logical drive creation encoder logic to use the actual physical drive block size for the array instead of a hard coded value.
  - Risk: Low
- Fixed an issue where the Drive.Identifiers.DurableName value for NVMe® drives did not conform to the standards regular expression.
  - Root cause: The code to separate DurableName for NVMe drives with a colon was not implemented.
  - Fix: Modified the DurableName string for NVMe drives to separate each pair of characters with a colon.
  - Risk: Low
- Fixed an issue where Volume CREATE requests succeeded with malformed data drive @odata.id strings.

- Root Cause: The validation of the `@odata.id` string did not enforce that the string contained at least one instance of the substring `"/Drives/"`.
- Fix: Modified the `@odata.id` string validation to enforce inclusion of the `"/Drives/"` substring prior to the drive collection index number.
- Risk: Low
- Fixed an issue where the `Volume.Identifiers.DurableName` value does not conform to the standards regular expression.
  - Root Cause: The code to appropriately format the `DurableName` string was not implemented.
  - Fix: Added hyphens to the `DurableName` string for Volumes as needed to meet the specification requirement.
  - Risk: Low
- Certain controller temperature sensor numeric sensors have had their `EntityType` changed from "I/O Controller" to "Add-in card".
- RDE READ on a Drive resource will now exclude the `Vendor` from `Drive.Name` property on some controllers.
- The Type 5 commands `QueryDownstreamDevices`, `GetDownstreamFirmwareParameters`, and `QueryDownstreamIdentifiers` will now report information for physical drives.
- The `Volume.WriteCachePolicy` property in a Redfish GET response for a Volume resource will have a value of `ProtectedWriteBack` when the controller battery is removed or goes missing.
- Updated all resource schema dictionaries to the latest version available in the 2021.4 schema bundle.
- On controllers that support managed SED encryption:
  - Redfish GET responses for a self-encrypting drive resource will publish the following `EncryptionStatus` values:
    - Unencrypted
    - Locked
    - Unlocked
    - Foreign
  - `Drive.Status.State` for a self-encrypting drive (SED) resource will be set to `StandbyOffline` in the following conditions:
    - SED is Foreign
    - SED is Locked (only for controller owned SEDs)
    - SED is controller owned and controller is waiting on SED adapter password
  - `Encrypted` property will be set to `True` on Redfish GET responses for Volume resources, which are secured using SED-based encryption.

## 2.3 Limitations

### 2.3.1 Firmware Limitations

#### 2.3.1.1 Limitations for Firmware Release 5.32 B0

This release includes the following firmware limitations:

- Deleting a secure SED volume using an older firmware that does not support the SED Local Key Management (LKM) feature can cause the physical drive status to be incorrect when moved back to SED LKM aware firmware.
  - Workaround: Delete the secure SED volume prior to down revving the firmware.

- A firmware update causes the UART log buffer (Serial Output Buffer) to be reinitialized, since the DDR gets reinitialized.
  - Workaround: None
- SATA drives attached to a non-Microchip expander may get into a failed state when upgrading the controller firmware from previous releases to this release due to the expander not clearing STP affiliation.
  - Workaround: Power cycle the expanders to clear the STP affiliation.
- A rare corner-case scenario where controller may hang during expander firmware update on multi-level expander/SEP device topology along with I/Os.
  - Workaround: After the enclosure firmware update, avoid enclosure Reset. It is recommended to download the new firmware and perform manual power cycle. This issue is intermittent and can cause a hang that requires a system reboot.
 

**Note:** This issue was mostly seen when using Linux OS.
- Controller cache will not be converted into 100% read cache, if any backup power source cable error, charge or charge timeout error occurs when expansion or transformation task is active.
  - Workaround: None
- Firmware downgrade is blocked if disk-based transformation is in-progress.
  - Workaround: Wait for the transformation to complete and retry the firmware downgrade.
- Transformation is blocked if,
  - Reboot after the firmware update is pending.
  - Flashed new firmware version is older than 5.32 B0.
    - Workaround: Reboot the system.
- Logical drive is not detected if,
  - Disk-based transformation is in-progress during logical drive movement to a different controller and the different controller has a firmware version older than 5.32 B0.
  - Firmware downgrade occurred while internal-cache based transformation was in progress, but the Backup Power Source failed before firmware activation.
    - Workaround: Move the logical drive to a controller with firmware version 5.32 B0 or later.
- Logical drives containing multiple predictive failed physical drives may be placed offline by the OS.
 

A logical drive with multiple predictive failed drives may see an I/O latency issue. After all but one of the predictive drives are replaced and the logical drive rebuild is completed, the I/O latency issue might still remain which can result in the OS marking the logical drive offline.

  - Workaround: Reboot the server to correct the I/O latency issue and get the logical drive back online.
- When Predictive Spare Rebuild is enabled, the logical drive metadata is not cleared on Predictive Failed drives after the host issues a “clear configuration” for the controller.
  - Workaround: Reboot the server and re-issue the clear configuration command to delete the RAID metadata from the predictive failed drives.
- PBSI shows dedicated hotspare type when auto replace hotspare is configured.
  - Workaround: None

### 2.3.1.2 Limitations for Firmware Release 1.32 Build 0

- Firmware release 1.32b0 may become unresponsive while attempting to flash firmware or execute other RAID logical drive operations.
  - Description: Refer to entry "Fixed an issue where firmware may become unresponsive while attempting to flash firmware or execute other RAID logical drive operations" in the Firmware fixes section.

- A fix for this issue is available in the 1.60 B0 firmware release. If a firmware flash failure is occurring, try the following workarounds:
  - Workaround: If there are no target devices (expanders or drives) attached to the controller, attach a target device to the controller and try the host management operation again.
  - Workaround: If the system is operating using UEFI, the HII tool can be used to flash the firmware to this release as outlined in the *Microchip SmartIOC 2100/SmartROC 3100 Installation and User's Guide (ESC-2170577)*, appendix entry "Updating the SmartIOC 2100/SmartROC 3100 Controller Firmware".
  - Workaround: If there are target devices attached to the controller and this issue occurs or none of the workarounds can be used, contact Microchip Support.

## 2.3.2 UEFI Limitations

### 2.3.2.1 Limitations for UEFI Build 2.2.4/Legacy BIOS Build 2.2.2

This release includes the following limitation:

- The Revert with PSID operation fails if the PSID for the SED is less than 32 characters (bytes).
  - Workaround: Run the following command in SEDUTIL CLI to revert the SED:

```
sedutil-cli --yesIreallywanttoERASEALLmydatausingthePSID <PSID> <device>
```

- In UEFI HII, when enabling the SED Local Key Management, the optional Controller Password feature is not supported and if enabled will result in the controller returning a failure to enable the SED Local Key Management.
  - Workaround: Enable the SED Local Key Management without enabling the Controller Password feature.

## 2.3.3 Driver Limitations

### 2.3.3.1 Limitations for Linux Driver Build 2.1.18-045

This release has the following Linux limitation:

- On AMD/RHEL 7.9 systems, the system might panic due to a bug in the IOMMU module. For details, see <https://lore.kernel.org>
  - Workaround: Disable the IOMMU setting option in BIOS.
- The following are the limitations of Multi-Actuator:
  - Supports only
    - HBA drive
    - Direct-Attached
    - Windows/Linux/VMware
    - Intel/AMD
    - UEFI mode (for multi-LUN display)
  - No Storage Manager support
  - No boot support
- On AMD/UEK6 systems, the system might hang during kdump if IOMMU is enabled.
  - Workaround: Disable IOMMU setting option in BIOS.
- RHEL driver injection (DUD) install where OS ISO is mounted as virtual media on BMC based servers (non-ILO). Installer will hang after driver injection. Reported on RHEL 8.5, 8.6 and 9.0.
  - Workaround: Load OS from USB device instead of virtual media.

Load OS from virtual media but initiation ISO verification (media test) during install followed by ESC to cancel media test.

- This release includes the following limitation when doing a driver injection (DUD) install. On some distributions (RHEL7.9, RHEL8.2, RHEL8.3, SLES15SP2, SLES15SP3), the DUD install will hang if an attached drive (either HBA mode or Logical Volume) has Write Cache enabled.
  - Workaround: There are two work-arounds for this issue:
    - Make sure the Write Cache is disabled for any attached drive.
    - For RHEL7.9/8.2/8.3, add `rd.driver.blacklist=smartpqi` to the grub entry along with `inst.dd`.
- Depending on hardware configurations, the `smartpqi expose_ld_first` parameter may not always work consistently.
  - Workaround: None
- When multiple controllers are in a system, `udev(systemd)` can timeout during `kdump/kexec` resulting in an incomplete `kdump` operation. The usual indication of the timeout is the console log entry: `"scsi_hostX: error handler thread failed to spawn, error = -4"`.
  - Workaround: Extend the `udev(systemd)` timeout during a `kdump` operation. The steps to increase the timeout for `udev(systemd)` are:
    - `vi /etc/sysconfig/kdump`
    - add `udev.event-timeout=300` to `KDUMP_COMMANDLINE_APPEND`
    - `systemctl restart kdump`
    - `systemctl status kdump`

### 2.3.3.2 Limitations for Windows Driver Build 1010.42.0.1020

This release includes the following limitation:

- The following are the limitations of Multi-Actuator:
  - Supports only
    - HBA drive
    - Direct-Attached
    - Windows/Linux/VMware
    - Intel/AMD
    - UEFI mode (for multi-LUN display)
  - No Storage Manager support
  - No boot support

### 2.3.3.3 Limitations for FreeBSD Driver Build 4280.0.1007

There are no known limitations for this release.

### 2.3.3.4 Limitations for Solaris Driver Build 11.4120.0.1005

There are no known limitations for this release.

### 2.3.3.5 Limitations for VMware Driver Build 4330.0.116

This release includes the following limitations:

- If the controller SED Encryption feature is "On" and locked, Datastores created from secured logical drives on the controller are not automatically mounted even after unlocking the controller, they are not visible through the ESXi hypervisor client.
  - Workaround: Use the command `vmkfstool -V or esxcli storage filesystem rescan`. Alternatively, you can also use the **Rescan** option from the **Devices** tab in the Hypervisor's Storage section. Any of these options solve the issue by forcing a rescan, causing the datastore to mount.

- The following are the limitations of Multi-Actuator:
  - Supports only
    - HBA drive
    - Direct-Attached
    - Windows/Linux/VMware
    - Intel/AMD
    - UEFI mode (for multi-LUN display)
  - No Storage Manager support
  - No boot support

## 2.3.4 Management Software Limitations

### 2.3.4.1 Limitations for Arconf/maxView Build B25335

This release includes the following limitation:

- The Revert with PSID operation fails if the PSID for the SED is less than 32 characters (bytes).
  - Workaround: Run the following command in SEDUTIL CLI to revert the SED:

```
sedutil-cli --yesIreallywanttoERASEALLmydatausingthePSID <PSID> <device>
```

- When enabling the SED Local Key Management, the optional Controller Password feature is not supported and if enabled will result in the controller returning a failure to enable the SED Local Key Management.
  - Workaround: Enable the SED Local Key Management without enabling the Controller Password feature.

### 2.3.4.2 Limitations for PLDM Release 6.10.14.0

This release includes the following PLDM limitations:

- Action `Storage.ResetToDefault` with a `ResetType` of 'ResetAll' is not supported when the controller has logical drives that are encrypted.
  - Workaround: None
- RDE update on `Drive.WriteCacheEnabled` is unsupported for physical drives that are part of a logical drive.
  - Workaround: None

## 2.3.5 Hardware Limitations

This release includes the following hardware limitations:

- Two Wire Interface (TWI) address conflicts can cause system DDR memory to not be discovered.
  - Description: The SmartRAID 3100 and SmartHBA 2100 boards include two TWI targets on the host-facing SMBUS interface with the following slave addresses:
    - 0xA0 – Field Replaceable Unit (FRU) SEEPROM
    - 0xDE – PBSI (default)

According to the JEDEC specification, the default TWI addresses for the DDR SPD is 0xA0-0xAE (the spec uses 7 bit addressing which is 0x50-0x57). On platform system board designs with SMBUS wiring that has both PCIe slots and DDR slots shared on the same TWI bus, the TWI devices for the DDR and Smart controller are exposed to address conflicts which can result in the system memory not being discovered. The Smart controller PBSI interface defaults to a value of 0xDE (0x6F in 7-bit addressing) and is not a problem unless it is changed to an address that conflicts with the JEDEC defined values. The Smart controller FRU SEEPROM is hardwired to 0xA0.



- Workaround: None available. If this issue is encountered, contact your Microchip support engineer to determine the next steps for your system.
- Performance with workaround: Not applicable
- Performance without workaround: Not applicable

### 3. Updating the Controller Firmware

This section describes how to update the board's firmware components to the latest release.

#### 3.1 Updating the Controller Firmware

This procedure describes how to prepare your board to be programmed with the latest firmware.

**Notes:**

1. If the running firmware is older than 1.98 and a transformation is in progress, complete the transformation before proceeding with the following steps to upgrade the firmware.
2. Complete these procedures exactly as described for proper functionality. If you do not follow all of the steps correctly, you could encounter unusual runtime behavior.

**Flashing the board to the latest firmware:**

This section describes how to update all the firmware components on Adaptec controller boards to the latest release.

**If the controller is currently running 1.60 b0 firmware or newer, follow these steps:**

1. **Mandatory:** Flash the target with the provided "SmartFWx100.bin" image with arconf/maxView software.
2. **Mandatory:** Use the OS shutdown/restart operation to gracefully reboot the system to complete the firmware update process.

**Note:**

After completing the firmware update, if the firmware version is still showing the prior version, retry the firmware update steps.

**If the controller is currently running 1.32 b0 firmware, follow these steps:**

1. **Mandatory:** Flash the target with the provided "SmartFWx100.bin" image with arconf/maxView software.
  - If the arconf/maxView software becomes unresponsive or hangs then power cycle the system to recover and refer to firmware limitation section [2.3.1.2. Limitations for Firmware Release 1.32 Build 0](#).
2. **Mandatory:** If flashing completes, use the OS shutdown/restart operation to gracefully reboot the system to complete the firmware update process.

**Note:**

After completing the firmware update, if the firmware version is still showing the prior version, retry the firmware update steps.

**If the controller is currently running 1.04 b0 firmware, follow these steps:**

1. **Mandatory:** Flash the controller with the provided "SmartFWx100\_v1.29\_b314.bin" image with arconf/maxView software.
2. **Mandatory:** Reboot the system to refresh all components.
3. **Mandatory:** Flash the target with the provided "SmartFWx100.bin" image with arconf/maxView software.
4. **Mandatory:** Use the OS shutdown/restart operation to gracefully reboot the system to complete the firmware update process.

At this point, the controller would be updated and would be ready to use. Install the SmartPQI driver and the latest version of the Arconf/maxView management utility to monitor and configure the controller.

**Note:** Downgrading firmware could lead to unexpected behavior due to an incompatibility in SEEPROMs between this release and the prior release.

## 4. Installing the Drivers

See the “*Microchip Adaptec® SmartRAID 3100 Series and SmartHBA 2100 Series Host Bus Adapters Installation and User's Guide (DS00004439B, previously ESC-2171547)*” for complete driver installation instructions.

## 5. Revision History

The revision history describes the changes that were implemented in the document. The changes are listed by revision, starting with the most current publication.

Revision	Date	Description
J	10/2023	SR 2.7.0 Patch Release with maxView™ version B25339.
H	07/2023	SR 2.8.0 Production Release
G	03/2023	SR 2.7.4 Production Release
F	11/2022	SR 2.7.2 Production Release
E	08/2022	SR 2.7.0 Production Release
D	03/2022	VMware driver version updated from 4250.0.120 to 4252.0.103
C	02/2022	SR 2.6.6 Production Release
B	12/2021	SR 2.6.4.1 Patch Release with maxView version B24713. Updated Fixes and Enhancements for maxView Storage Manager/ARCCONF section for log4j vulnerabilities.
A	11/2021	SR 2.6.4 Production Release with firmware version 4.72 B0 (Previously ESC-2161026)
29	04/2021	SR 2.6.2 with firmware version 4.11 B0
28	04/2021	SR 2.6.1.1 with VMware driver version 4054.2.118.
27	03/2021	SR 2.6.1 with VMware driver version 4054.1.103.
26	02/2021	SR 2.6 Production Release
25	10/2020	SR 2.5.4 Production Release
24	08/2020	SR 2.5.2.2 Production Release with Firmware 3.00
23	03/2020	SR 2.5.2 Production Release with Firmware 2.93
22	03/2020	SR 2.5 Production Release with Firmware 2.66
21	02/2020	SR 2.5.2 Production Release
20	10/2019	SR 2.5 Production Release
19	09/2019	Updated for SR 2.4.8.1 (fw v2.31 Build 0)
18	08/2019	Updated for SR 2.4.8
17	01/2019	SR2.4 Production Release
16	06/2018	SR2.3 Production Release
15	06/2018	Updated for RC Release
14	10/2017	Update supported OSs
13	10/2017	First Production Release
1-12	06/2016 to 07/2017	Pre-Production Release.

## Microchip Information

### The Microchip Website

Microchip provides online support via our website at [www.microchip.com/](http://www.microchip.com/). This website is used to make files and information easily available to customers. Some of the content available includes:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user’s guides and hardware support documents, latest software releases and archived software
- **General Technical Support** – Frequently Asked Questions (FAQs), technical support requests, online discussion groups, Microchip design partner program member listing
- **Business of Microchip** – Product selector and ordering guides, latest Microchip press releases, listing of seminars and events, listings of Microchip sales offices, distributors and factory representatives

### Product Change Notification Service

Microchip’s product change notification service helps keep customers current on Microchip products. Subscribers will receive email notification whenever there are changes, updates, revisions or errata related to a specified product family or development tool of interest.

To register, go to [www.microchip.com/pcn](http://www.microchip.com/pcn) and follow the registration instructions.

### Customer Support

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Embedded Solutions Engineer (ESE)
- Technical Support

Customers should contact their distributor, representative or ESE for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in this document.

Technical support is available through the website at: [www.microchip.com/support](http://www.microchip.com/support)

### Microchip Devices Code Protection Feature

Note the following details of the code protection feature on Microchip products:

- Microchip products meet the specifications contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is secure when used in the intended manner, within operating specifications, and under normal conditions.
- Microchip values and aggressively protects its intellectual property rights. Attempts to breach the code protection features of Microchip product is strictly prohibited and may violate the Digital Millennium Copyright Act.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of its code. Code protection does not mean that we are guaranteeing the product is “unbreakable”. Code protection is constantly evolving. Microchip is committed to continuously improving the code protection features of our products.

### Legal Notice

This publication and the information herein may be used only with Microchip products, including to design, test, and integrate Microchip products with your application. Use of this information in any other manner violates these terms. Information regarding device applications is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure

that your application meets with your specifications. Contact your local Microchip sales office for additional support or, obtain additional support at [www.microchip.com/en-us/support/design-help/client-support-services](http://www.microchip.com/en-us/support/design-help/client-support-services).

THIS INFORMATION IS PROVIDED BY MICROCHIP "AS IS". MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE, OR WARRANTIES RELATED TO ITS CONDITION, QUALITY, OR PERFORMANCE.

IN NO EVENT WILL MICROCHIP BE LIABLE FOR ANY INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, OR CONSEQUENTIAL LOSS, DAMAGE, COST, OR EXPENSE OF ANY KIND WHATSOEVER RELATED TO THE INFORMATION OR ITS USE, HOWEVER CAUSED, EVEN IF MICROCHIP HAS BEEN ADVISED OF THE POSSIBILITY OR THE DAMAGES ARE FORESEEABLE. TO THE FULLEST EXTENT ALLOWED BY LAW, MICROCHIP'S TOTAL LIABILITY ON ALL CLAIMS IN ANY WAY RELATED TO THE INFORMATION OR ITS USE WILL NOT EXCEED THE AMOUNT OF FEES, IF ANY, THAT YOU HAVE PAID DIRECTLY TO MICROCHIP FOR THE INFORMATION.

Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

## Trademarks

The Microchip name and logo, the Microchip logo, Adaptec, AVR, AVR logo, AVR Freaks, BesTime, BitCloud, CryptoMemory, CryptoRF, dsPIC, flexPWR, HELDO, IGLOO, JukeBlox, KeeLoq, Kleer, LANCheck, LinkMD, maXStylus, maXTouch, MediaLB, megaAVR, Microsemi, Microsemi logo, MOST, MOST logo, MPLAB, OptoLyzer, PIC, picoPower, PICSTART, PIC32 logo, PolarFire, Prochip Designer, QTouch, SAM-BA, SenGenuity, SpyNIC, SST, SST Logo, SuperFlash, Symmetricom, SyncServer, Tachyon, TimeSource, tinyAVR, UNI/O, Vectron, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

AgileSwitch, APT, ClockWorks, The Embedded Control Solutions Company, EtherSynch, Flashtec, Hyper Speed Control, HyperLight Load, Libero, motorBench, mTouch, Powermite 3, Precision Edge, ProASIC, ProASIC Plus, ProASIC Plus logo, Quiet-Wire, SmartFusion, SyncWorld, Temux, TimeCesium, TimeHub, TimePictra, TimeProvider, TrueTime, and ZL are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, Augmented Switching, BlueSky, BodyCom, Clockstudio, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, Espresso T1S, EtherGREEN, GridTime, IdealBridge, In-Circuit Serial Programming, ICSP, INICnet, Intelligent Paralleling, IntelliMOS, Inter-Chip Connectivity, JitterBlocker, Knob-on-Display, KoD, maxCrypto, maxView, memBrain, Mindi, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICkit, PICtail, PowerSmart, PureSilicon, QMatrix, REAL ICE, Ripple Blocker, RTAX, RTG4, SAM-ICE, Serial Quad I/O, simpleMAP, SimpliPHY, SmartBuffer, SmartHLS, SMART-I.S., storClad, SQL, SuperSwitcher, SuperSwitcher II, Switchtec, SynchroPHY, Total Endurance, Trusted Time, TSHARC, USBCheck, VariSense, VectorBlox, VeriPHY, ViewSpan, WiperLock, XpressConnect, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

The Adaptec logo, Frequency on Demand, Silicon Storage Technology, and Symmcom are registered trademarks of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.

© 2023, Microchip Technology Incorporated and its subsidiaries. All Rights Reserved.

ISBN: 978-1-6683-3361-7

## **Quality Management System**

For information regarding Microchip's Quality Management Systems, please visit [www.microchip.com/quality](http://www.microchip.com/quality).



# Worldwide Sales and Service

AMERICAS	ASIA/PACIFIC	ASIA/PACIFIC	EUROPE
<p><b>Corporate Office</b> 2355 West Chandler Blvd. Chandler, AZ 85224-6199 Tel: 480-792-7200 Fax: 480-792-7277 Technical Support: <a href="http://www.microchip.com/support">www.microchip.com/support</a> Web Address: <a href="http://www.microchip.com">www.microchip.com</a></p> <p><b>Atlanta</b> Duluth, GA Tel: 678-957-9614 Fax: 678-957-1455</p> <p><b>Austin, TX</b> Tel: 512-257-3370</p> <p><b>Boston</b> Westborough, MA Tel: 774-760-0087 Fax: 774-760-0088</p> <p><b>Chicago</b> Itasca, IL Tel: 630-285-0071 Fax: 630-285-0075</p> <p><b>Dallas</b> Addison, TX Tel: 972-818-7423 Fax: 972-818-2924</p> <p><b>Detroit</b> Novi, MI Tel: 248-848-4000</p> <p><b>Houston, TX</b> Tel: 281-894-5983</p> <p><b>Indianapolis</b> Noblesville, IN Tel: 317-773-8323 Fax: 317-773-5453 Tel: 317-536-2380</p> <p><b>Los Angeles</b> Mission Viejo, CA Tel: 949-462-9523 Fax: 949-462-9608 Tel: 951-273-7800</p> <p><b>Raleigh, NC</b> Tel: 919-844-7510</p> <p><b>New York, NY</b> Tel: 631-435-6000</p> <p><b>San Jose, CA</b> Tel: 408-735-9110 Tel: 408-436-4270</p> <p><b>Canada - Toronto</b> Tel: 905-695-1980 Fax: 905-695-2078</p>	<p><b>Australia - Sydney</b> Tel: 61-2-9868-6733</p> <p><b>China - Beijing</b> Tel: 86-10-8569-7000</p> <p><b>China - Chengdu</b> Tel: 86-28-8665-5511</p> <p><b>China - Chongqing</b> Tel: 86-23-8980-9588</p> <p><b>China - Dongguan</b> Tel: 86-769-8702-9880</p> <p><b>China - Guangzhou</b> Tel: 86-20-8755-8029</p> <p><b>China - Hangzhou</b> Tel: 86-571-8792-8115</p> <p><b>China - Hong Kong SAR</b> Tel: 852-2943-5100</p> <p><b>China - Nanjing</b> Tel: 86-25-8473-2460</p> <p><b>China - Qingdao</b> Tel: 86-532-8502-7355</p> <p><b>China - Shanghai</b> Tel: 86-21-3326-8000</p> <p><b>China - Shenyang</b> Tel: 86-24-2334-2829</p> <p><b>China - Shenzhen</b> Tel: 86-755-8864-2200</p> <p><b>China - Suzhou</b> Tel: 86-186-6233-1526</p> <p><b>China - Wuhan</b> Tel: 86-27-5980-5300</p> <p><b>China - Xian</b> Tel: 86-29-8833-7252</p> <p><b>China - Xiamen</b> Tel: 86-592-2388138</p> <p><b>China - Zhuhai</b> Tel: 86-756-3210040</p>	<p><b>India - Bangalore</b> Tel: 91-80-3090-4444</p> <p><b>India - New Delhi</b> Tel: 91-11-4160-8631</p> <p><b>India - Pune</b> Tel: 91-20-4121-0141</p> <p><b>Japan - Osaka</b> Tel: 81-6-6152-7160</p> <p><b>Japan - Tokyo</b> Tel: 81-3-6880-3770</p> <p><b>Korea - Daegu</b> Tel: 82-53-744-4301</p> <p><b>Korea - Seoul</b> Tel: 82-2-554-7200</p> <p><b>Malaysia - Kuala Lumpur</b> Tel: 60-3-7651-7906</p> <p><b>Malaysia - Penang</b> Tel: 60-4-227-8870</p> <p><b>Philippines - Manila</b> Tel: 63-2-634-9065</p> <p><b>Singapore</b> Tel: 65-6334-8870</p> <p><b>Taiwan - Hsin Chu</b> Tel: 886-3-577-8366</p> <p><b>Taiwan - Kaohsiung</b> Tel: 886-7-213-7830</p> <p><b>Taiwan - Taipei</b> Tel: 886-2-2508-8600</p> <p><b>Thailand - Bangkok</b> Tel: 66-2-694-1351</p> <p><b>Vietnam - Ho Chi Minh</b> Tel: 84-28-5448-2100</p>	<p><b>Austria - Wels</b> Tel: 43-7242-2244-39 Fax: 43-7242-2244-393</p> <p><b>Denmark - Copenhagen</b> Tel: 45-4485-5910 Fax: 45-4485-2829</p> <p><b>Finland - Espoo</b> Tel: 358-9-4520-820</p> <p><b>France - Paris</b> Tel: 33-1-69-53-63-20 Fax: 33-1-69-30-90-79</p> <p><b>Germany - Garching</b> Tel: 49-8931-9700</p> <p><b>Germany - Haan</b> Tel: 49-2129-3766400</p> <p><b>Germany - Heilbronn</b> Tel: 49-7131-72400</p> <p><b>Germany - Karlsruhe</b> Tel: 49-721-625370</p> <p><b>Germany - Munich</b> Tel: 49-89-627-144-0 Fax: 49-89-627-144-44</p> <p><b>Germany - Rosenheim</b> Tel: 49-8031-354-560</p> <p><b>Israel - Ra'anana</b> Tel: 972-9-744-7705</p> <p><b>Italy - Milan</b> Tel: 39-0331-742611 Fax: 39-0331-466781</p> <p><b>Italy - Padova</b> Tel: 39-049-7625286</p> <p><b>Netherlands - Drunen</b> Tel: 31-416-690399 Fax: 31-416-690340</p> <p><b>Norway - Trondheim</b> Tel: 47-72884388</p> <p><b>Poland - Warsaw</b> Tel: 48-22-3325737</p> <p><b>Romania - Bucharest</b> Tel: 40-21-407-87-50</p> <p><b>Spain - Madrid</b> Tel: 34-91-708-08-90 Fax: 34-91-708-08-91</p> <p><b>Sweden - Gothenberg</b> Tel: 46-31-704-60-40</p> <p><b>Sweden - Stockholm</b> Tel: 46-8-5090-4654</p> <p><b>UK - Wokingham</b> Tel: 44-118-921-5800 Fax: 44-118-921-5820</p>